



Kubernetes
Forum *Seoul*

Re-architecting Data Platform with Kubernetes

SeungYong Oh, Devsisters Corp.

오승용, 데브시스터즈



DEVSISTERS

- Mobile Game Publisher: including **Cookie Run Series**
- 10+ Million App Downloads, 2M+ Monthly Active Users

"Creating the Best Player Experiences through Excellent Content, Service, and Technology!"

"탁월한 기술, 서비스, 콘텐츠로 전 세계 고객에게 최고의 경험을 선사합니다."

SeungYong Oh (오 승 용)

- DevOps & Data Engineer at Devsisters
 - Also worked as a game server developer for Cookie Run series
 - Kubernetes user from 2016
 - Adopted Kubernetes for development/testing environment for Cookie Run OvenBreak
 - NDC 2017: Kubernetes로 개발서버 간단히 찍어내기
- <https://www.slideshare.net/seungyongoh3/ndc17-kubernetes>

- 2014~2015: Proof-of-Concept level evaluation
- 2016: **Dev/testing environment game infra** for **Cookie Run: OvenBreak**
- 2017~: Dev/testing environment **game infra platform**
- 2019~: **Production-level game infra platform**
 - Cookie Run for Kakao, Hello Brave Cookies
- 2019~: **Data Platform** on Kubernetes
 - **Still ongoing**

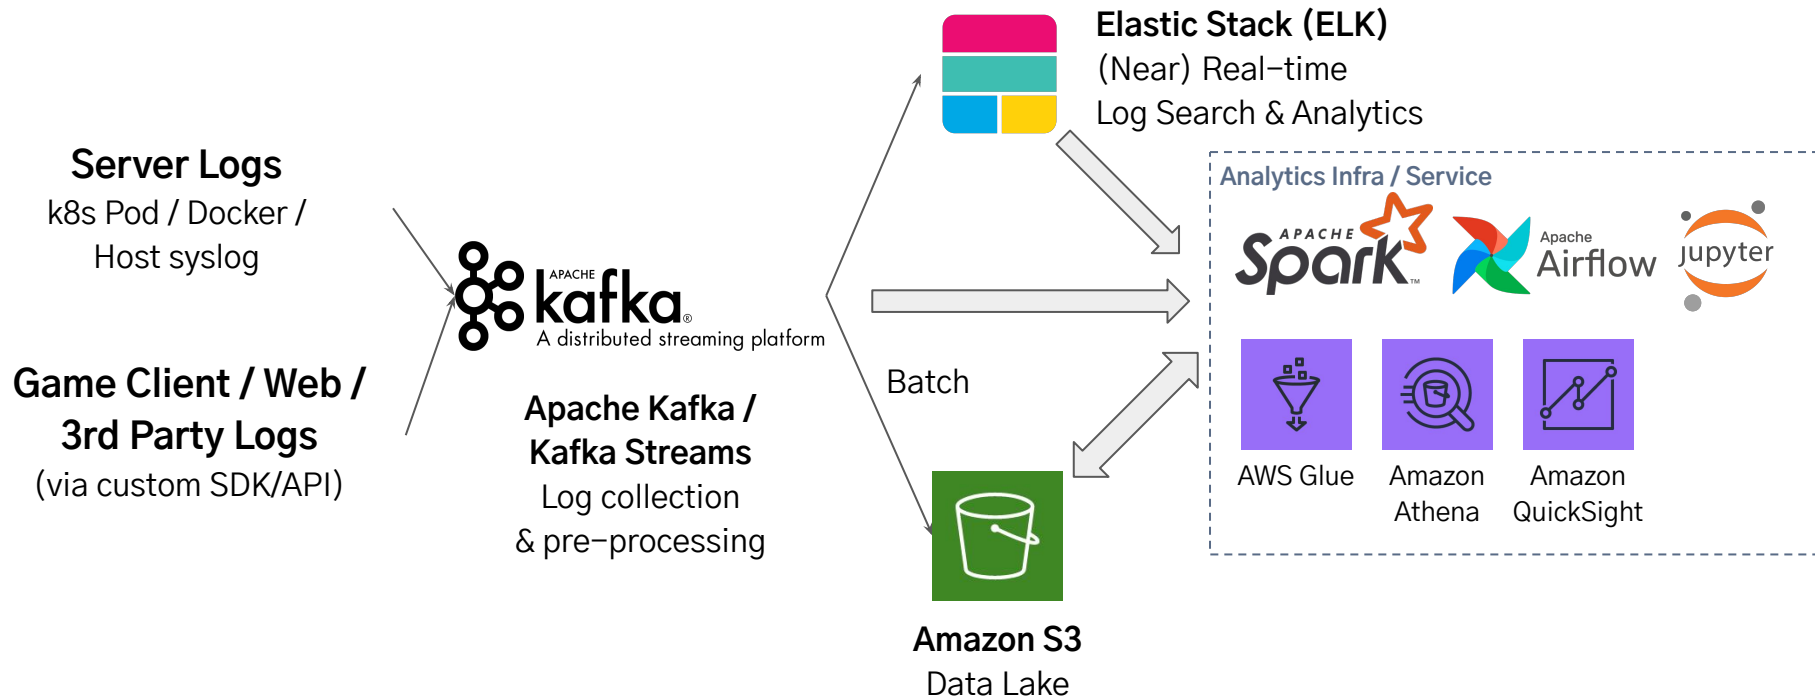
- Central logging infra with various data sources

Using those logs,

- Search
 - Search server/client logs for debug or operational purpose
 - Search specific user's action logs for customer service purpose
- Applications
 - A/B Testing, Machine Learning, ..
- And of course, Analytics!

- Analytics: Analytics for Everyone!
 - Everyone?
 - Not only limited to data scientist, analyst
 - CEO, Project Manager, Marketer, Game Designer, ...
 - KPI(Key Performance Indicator) Service: Active users, Revenue, Retention,...
 - Ad-hoc query/analytics environment
 - Fully programmable environment for Data Scientist / Engineer
 - SQL query / clickable interface for many users

Data Platform @ Devsisters



\$SOME_PROJECT_NAME supports K8S!



Many Big data & analytics related projects support Kubernetes in 2019!

- Hadoop, Spark, Airflow, Kubeflow, JupyterHub, Hue, Zeppelin, Superset, Kafka, Elastic Stack, etc.

Support?

- Case 1: (Official/Popular) Helm Chart
- Case 2: K8S Operator for the project
- Case 3: Cloud Native / Kubernetes-specific Integrations

But, does it mean we should move to K8S?



Many projects are stateful apps, especially in Big data / analytics

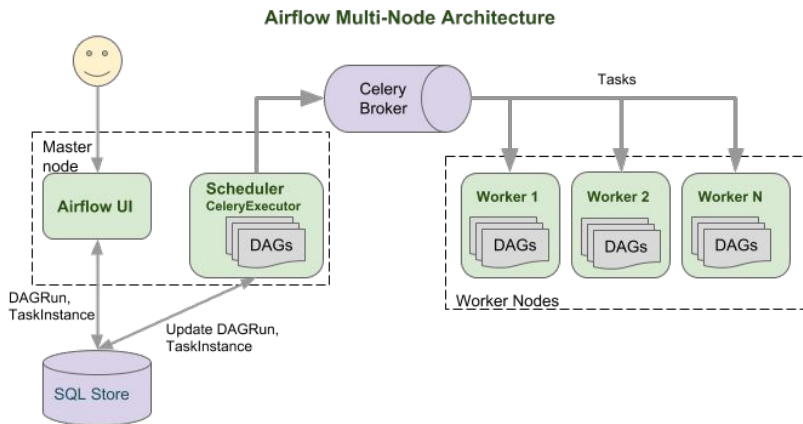
- Considerable risks / maintenance costs
 - Node Upgrade? Retirement? Reallocation is **NOT** easy
- Benefits?
 - Bin-packing? 🤔 (especially in public cloud)
 - Autoscaling? 🤔
 - Scheduling? 🤔
 - Already there are some advanced scheduler/resource manager (ex. Hadoop YARN)

Actually, there are some (great!) benefits!

Benefit #1: Deployment / Ops Made Easy

Example: Apache Airflow

Apache Airflow: Workflow Management tool



Quick Deployment with

- helm install stable/airflow

or

- Airflow operator

<https://github.com/GoogleCloudPlatform/airflow-operator/>

Source:

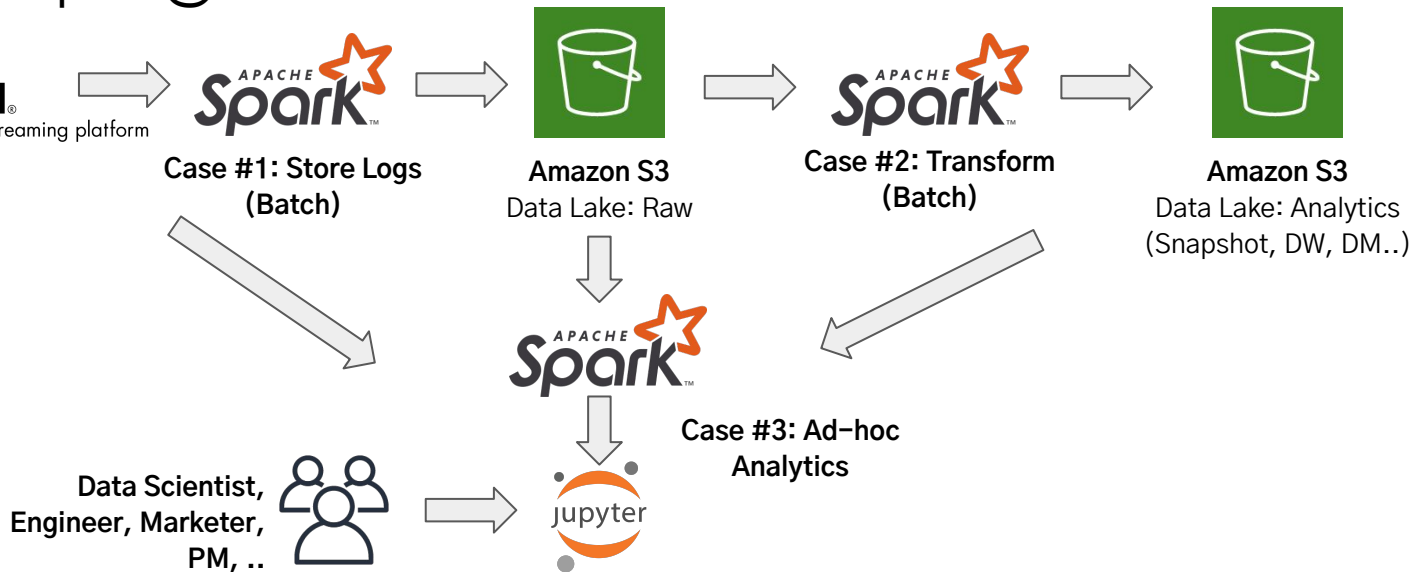
<https://github.com/GoogleCloudPlatform/airflow-operator/blob/master/docs/design.md>

Benefit #2: Analytics - Easy and Efficient

Apache Spark @ Devsisters

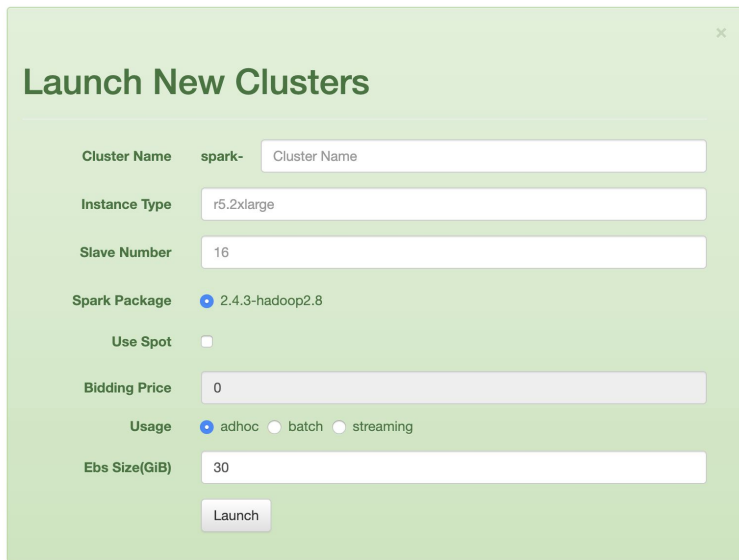
A unified analytics engine for large-scale data processing

Use-cases of Spark @ Devsisters:



Apache Spark @ Devsisters

Self-service web interface for Spark Cluster Provision



Launch New Clusters

Cluster Name: spark-

Instance Type:

Slave Number:

Spark Package: ☒ 2.4.3-hadoop2.8

Use Spot: ☐


Bidding Price:

Usage: ☒ adhoc ☐ batch ☐ streaming

Ebs Size(GiB):

Launches Spark Cluster (standalone mode)

- 1 master node with Jupyter Notebook
- multiple worker nodes

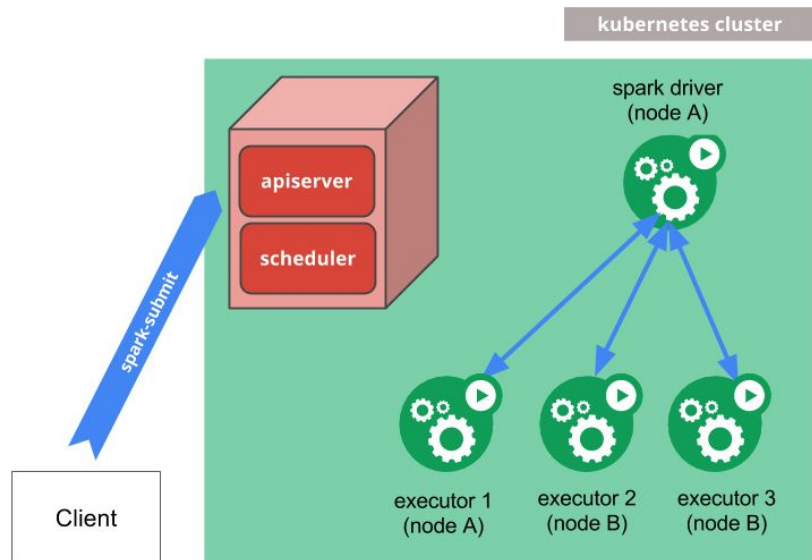
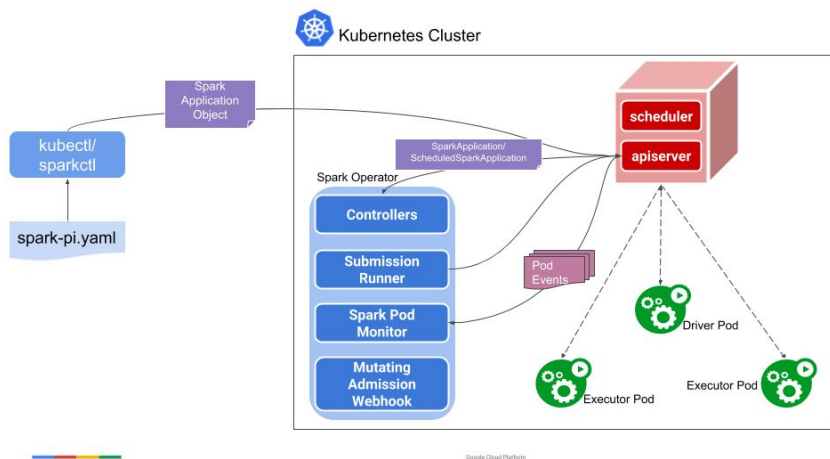


So, anyone can have their own Spark Cluster with one-click! Looks Good?

- Many users (ex. Marketing Manager) only want to focus on data analysis!
 - r5.xlarge?
 - spot? bidding price?
 - Waiting few minutes for spark provisioning, with some failures
 - Should backup their files before terminating cluster
- Cost ineffective
 - idle resources
 - forget to terminate cluster after use

Then sharing a centralized cluster? NO! (access control issues etc.)

Spark on Kubernetes



K8S Operator for Apache Spark

<https://github.com/GoogleCloudPlatform/spark-on-k8s-operator>

Apache Spark - Native K8S Scheduler (experimental)

<https://spark.apache.org/docs/latest/running-on-kubernetes.html>

Jupyter + Spark on K8S



```
# Execute below in Jupyter Pod
PYSPARK_DRIVER_PYTHON=jupyter \
PYSPARK_DRIVER_PYTHON_OPTS='notebook --allow-root' \
./bin/pyspark \
  --master k8s://https://$KUBERNETES_SERVICE_HOST:$KUBERNETES_PORT_443_TCP_PORT/ \
  --name spark-$APP_NAME \
  --conf spark.executor.instances=$EXECUTOR_COUNT \
  --conf spark.kubernetes.container.image=<<custom built spark container>> \
  --conf spark.kubernetes.namespace=spark-test \
  --conf spark.kubernetes.authenticate.driver.serviceAccountName=spark-test-server \
  --conf spark.driver.host=spark-driver \
  --conf spark.driver.port=5555 \
  --conf spark.driver.bindAddress=0.0.0.0 \
  --conf spark.driver.blockManager.port=5556 \
  --conf spark.kubernetes.executor.annotation.iam.amazonaws.com/role=<<AWS IAM role (KIAM)>>
```

- But users need much easier interface!



- Multi-user version of Jupyter Notebook, support k8s natively
- Launch a notebook pod per user
 - Data is persisted (stored in PersistentVolume)
- Let's Integrate Spark on K8S with JupyterHub !

JupyterHub + Spark on K8S



- Demo

Sign in with keycloak

Spawner Options

☐ **Jupyter notebook**
Jupyter notebook with spark

☒ **Jupyter lab**
Jupyter lab with spark

Custom environment variables

SPARK_EXECUTOR_COUNT

2

SPARK_EXECUTOR_MEMORY

1G

Spawn

FileEditViewRunKernelTabsSettingsHelp

TERMINAL SESSIONS

KERNEL SESSIONS

Untitled1.ipynbSHUT DOWN

Untitled1.ipynb

Python 3

[1]: spark

[1]: SparkSession - hive

SparkContext

Spark UI

Versionv2.4.4

Masterk8s://https://10.100.0.1:443/

AppNamespark-app-jupyterhub-kimtkyeom


[]:

01Python 3 | Idle

Mode: EditLn 1, Col 1Untitled1.ipynb

But many users only need SQL interface

- Spark has distributed SQL Engine
<https://spark.apache.org/docs/latest/sql-distributed-sql-engine.html>
 - Thrift JDBC/ODBC server (corresponds to the HiveServer2)
- Integration with Apache Superset (Business Intelligence web app)



Untitled Query



Database: sparksql spark_thrift_server

Schema: sdkdemo_analytics_dev

See table schema (7 in sdkdemo_analytics_dev)

Select table or type table name

topic_platform

topic_game

latest partition: type=track/partition

app	VARCHAR
context	VARCHAR
env	VARCHAR
event	VARCHAR
fgsid	VARCHAR
gamecode	VARCHAR
level	VARCHAR
logtype	VARCHAR
message	VARCHAR
messageid	VARCHAR
...	...

```
1 SELECT DISTINCT(mid) FROM analytics_topic_analytics
2 WHERE type='identify' AND partition='login' AND date = '2019-11-18'
```

Run Query

Save Query

Share Query

LIMIT 1000



Results

Query History

Preview: `topic_platform`

Preview: `topic_game`

Preview: `

Explore

.CSV

Clipboard

mid

SQRBT0891

DDRLH5303

GNSRY2100

For applications?

- Example #1:
 - Send push messages to users that played game more than a hour yesterday
 - JDBC!
- Example #2:
 - Give special items to users that has some possibility not to play the game anymore
 - spark-submit?

Benefit #3: Fine-tuning Access Controls

- Web Apps: Use Auth Proxy Sidecar
- For apps that don't support RBAC:
 - Pod-level data access control (S3 in Devsisters' case)
 - In AWS, with KIAM or AWS IAM Authenticator
 - Or Node-level (with nodeSelector)
 - Pods/Deployment per each group/roles + Benefit of Bin-packing

Then, Let's move to Kubernetes Now..?!

- Spark on Kubernetes
 - Doesn't support dynamic allocation / external shuffle service yet
 - You can't change count of executors
 - Pod Templating - available on 3.0 Preview
 - Even You can't specify service account of executor pod now
- Other Many Projects also in alpha stage,
which requires code-level modifications / forking
to meet the team's requirements
- If storage & computing resource is coupled, maybe little more hard to move

But it's worth to move!



- A Data Engineer who doesn't have prior knowledge of Kubernetes could implement Spark on K8s/JupyterHub/Superset customizations less than 1.5 months (expected more than 4 months for non-k8s env)
- Kubernetes' abstraction allows users to focus on what they can do best